Suitability of the Internet for Electric Power Transmission Realtime Information Networks[1]

(A white paper. Version 2.1.)

William E. Johnston
Ernest Orlando Lawrence Berkeley National Laboratory, University of California
(A United States Department of Energy Laboratory)
Berkeley, CA 94720

1.0 Abstract

The U. S. Federal Energy Regulatory Commission, in a Notice of Proposed Rulemaking ("NOPRA") (Docket No. RM95-8-000, March 29, 1995) has proposed a set of actions to move the Electric Power Transmission system to open competition. (See [Johnston-95].) The technical aspects of the actions involve defining what information needs to be made available, and how that information is to be made available. The "how" is to be through open electronic data networks called "Realtime Information Networks" (RINs). The Electric Power Research Institute (EPRI) has been tasked with organizing a set of workshops to collect input and present a White Paper to FERC that proposes an implementation of the RIN [EPRI-95a]. One of these workshops was held in Chicago on Sept. 7-8, 1995. In the course of this Workshop, it was proposed that the Internet be used as the mechanism to access RIN data. Several issues and questions arose concerning the use of the Internet for this purpose. This paper attempts to address some of these issues and questions. In particular, it addresses:

- An overview of the technical aspects of the Internet (what it is, how it works, what are its capabilities, and what it is used for);
- How one gains access to the Internet;
- Security issues in the Internet environment;
- The cost of using the Internet.

This paper provides information to support the thesis that the Internet is a mature, reliable, and widely used infrastructure for the exchange of scientific, engineering, and financial data, and provides secure interaction when the end-user systems and applications are properly designed and managed.

Additionally, this paper contains a strawman design for a RIN based on an informal set of requirements and desired functionality expressed verbally at the Chicago workshop. This design is intended to clarify some issues and provide a vehicle for further discussion. The relationship between EDI and the Internet is discussed briefly, and an annotated bibliography is provided.

Finally, many attendees of the Workshop felt that the name (Realtime Information Network) was too generic since there are many such networks in many different disciplines, and suggested changing the name to "Transmission Service-providers Information Network", or TSIN, and the name "TSIN" is used throughout this white paper.

---

1. This work is supported by the Director, Office of Energy Research, Office of Computation and Technology Research, Mathematical, Information, and Computational Sciences (MICS) Division, of the U. S. Department of Energy under Contract No. DE-AC03-76SF00098 with the University of California. This document is LBNL report number LBL-37767.

This paper may also be found at http://www-itg.lbl.gov/~johnston/EDM/

The author may be reached at wejohnston@lbl.gov, http://www-itg.lbl.gov, tel; 510-486-5014, fax: 510-486-6363.

2.0 Contents

3.0 The Internet

3.1 What is the Internet?

The term “Internet” is used in several ways:

> 1) “Internet” refers to the global Internet - the world’s largest packet switched data communications network
>
> 2) “internet” (spelled with a lower case “i”) refers to private (isolated) packet data networks that are based on Internet technology
>
> 3) “Internet technology” refers to the Internet Protocols (addressing and data transport), routing algorithms (how the path from source to destination is found), network management, tools, practices, etc., used in the Internet.

The Internet, by its original design, provides a sharable and survivable communications infrastructure. It is constructed of “routers” (packet switches) interconnected by many different “link-level” network technologies. In the wide area, these links are usually provided by the telecommunications industry, and include, for example:

- dedicated (e.g. “T-carrier”) circuits;
- Asynchronous Transfer Mode (ATM) over SONET (synchronous optical networks) (ATM is a new digital communications technology being deployed by the telecommunications industry. It is also being used for high speed LANs.)
- dial-up analogue and digital (ISDN) telephone circuits.

In local area networks (LANs of the size of offices, buildings, or campuses), Ethernet and FDDI are typical link-level technologies. The Internet IP protocols run over all of these different link-level networks, providing “inter-networking”, hence “internet protocol” or “IP”. Transport protocols like TCP, UDP, and RTP use the IP network protocol for addressing and packet-level operations. (See [Hedrick] and [Comer].)

The global Internet is a loose collection of hundreds of thousands of routers operated by tens of thousands of different organizations (both users and providers). It is a consensual, ad hoc, collaboration; there is no central authority and no formal development plan. The Internet is usually conceptualized as a large “cloud” of connectivity whose internal structure is dynamic and not visible to the communicating end-user systems.

Internet Service Providers (ISPs) are organizations (commercial and otherwise) that supply end-user system access to the cloud. (ISPs may, or may not, operate parts of the network “inside” the cloud.)

Access to the Internet by user systems is usually characterized by the bandwidth from the Internet communication infrastructure to the user site or computer system. This access may be by

- permanent connection (typical for a company or campus that connects their internal LAN to the Internet);
- a connection-on-demand (either user initiated, or by a router that dials automatically when traffic is present).

3.2 Internet Connectivity, Data Transport, and Content

The provision of connectivity and transport in the Internet using Internet Service Providers is analogous to the cable TV distribution system and the cable drops to your home provided by local cable TV systems. These systems (at least today) have the model of transport and connectivity being provided by one set of companies (cable operators) and content is supplied by another (e.g. TV Networks). On the Internet, transport and connectivity are supplied by Internet Service Providers, and content is provided by any connected computer system that chooses to do so.

Content takes many forms on the Internet:

Information (content) on the Internet takes the form of many combinations of:

- documents
- data
- interactive programs
- images
- audio
- video
- topical conversation (news groups, bulletin boards, e-mail lists)
- messages (e-mail), etc.

This information may be delivered:

- statically (retrieved by users from archives)
- dynamically (by autonomous servers, e.g. e-mail)
- in real-time - e.g.
  - Internet Relay Chat (a technology for interactive, on-line discussion - see [IRC])
  - Mbone multimedia teleconferencing [Mbone]
  - specialized timesharing systems such as Delphi, America Online, Prodigy, etc.

Many different Internet tools are used to search for and deliver content via the Internet (see [Liu-95]). For example:

- “ftp” retrieves files whose locations are known
- “archie” searches a database of file names and Internet site locations

- “gopher” searches a database of short descriptions of files
- The World Wide Web (WWW) provides a “live-link” hypertext system and full text indexing of all of the documents on the Web

3.3 Clients and Servers

The client-sever concept grew out of the ARPA funded, joint development of Unix and Internet technology in the early 1970s. Servers are applications that organize and provide access to many different kinds of content (data, directory lookup, e-mail, etc.). Clients request and retrieve the content from servers on behalf (typically) of human users (e.g. WWW browsers such as Netscape and Mosaic). Servers are always available to respond to requests, clients are typically created on-demand (e.g. when a human starts a program). Modern servers and clients are “network-based”: servers may run on any system connected to your network and clients anywhere in the network may access those servers. In the case of internet servers and clients the “network” may be the global Internet.

There are many types of servers and clients that operate automatically, both within computer systems and between Internet connected computer systems, to do all sorts of “housekeeping” functions (e.g. timekeeping, system directory services, printing, e-mail, user validation, etc.)

3.4 Protocols

Protocols are the rules that are used to specify how to do addressing, format data, communicate control requests, etc., over a data communication link.

Network-level protocols specify, in a very general way, how to transport data (e.g. the TCP protocol). (See [Hedrick] for a brief introduction to Internet protocols, and [Comer] for an in-depth coverage.)

Application-level protocols organize the communication of specific types of data and actions between clients and servers, for example (just to illustrate a few of the many application protocols):

- ftp (file transfer protocol - also the name of the client application, in this case)
- smtp (Internet mail)
- http (WWW protocol)
- ntp (the network time protocol)

3.5 What is the World Wide Web?

The World Wide Web has become the primary way to access much of the information available on the Internet. The Web consists of well over 100,000 sites (servers) that provide access to millions of interconnected “documents.” (Hence the term “web.”)

The basic way of representing information (or “content”) on the WWW is though what are called “hypertext” documents. Hypertext provides a construct called a “link” which, when activated by the user, causes a “jump” to some other document. One of the innovations that has made the WWW so powerful is the generalization of hypertext links that encompass not just being able to jump to another text document, but able to point to any sort of generalized “document” - for example, images, audio and video clips, and (with Java [Java]) programs that execute on your local system to perform some useful task, etc. A further generalization allows the link to invoke a program (via the “common gateway interface” (CGI) on the server side, where a program is invoked on behalf of the remote user) with input provided by the user through “forms” on the client (browser) side.

The combination of these extensions to the classical notion of hypertext allows people to build interactive program interfaces that are presented via WWW browsers such as Netscape and Mosaic and that provide much of the functionality of a window system[2], but in a platform and window system independent fashion. These visual and interactive WWW "documents" are now uniformly accessible from Microsoft Windows, Macintosh, and Unix platforms. For an example of a 3D, interactive, graphical visualization based on this technology, see http://www-itg.lbl.gov/ITG.hm.pg.docs/dissect/info.html [Frog]. The platform on which the server runs is irrelevant to the client because all of the interaction between the server and the client is governed by the hypertext transfer protocol ("http") application protocol that runs over the Internet.

There are currently many books on creating WWW sites and documents. (See, for example, [Liu-95] or [Lemay-95].)

4.0 How Does One Gain Access to the Internet?

End-user access to the Internet is provided by Internet Service Providers (ISPs). An ISP might be a University networking group, a corporate networking group, an independent commercial service provider, and (increasingly) telephone companies. (Sprint, MCI, and AT&T all provide Internet access service, as do many of the regional Bell operating companies.) Major players in the cable TV industry have also indicated that they will supply Internet access.

In all cases, individual end-systems (whether your lap-top PC or your Department file server) connect to Internet routers operated by service providers. The service provider connects that router to other Internet routers, or directly to one of the Internet "backbone" networks.

The end-user computer connects to the ISP routers at the "edges" of the Internet cloud through any of a variety of interfaces: Ethernet, dial-up modems, ISDN, etc. The service provider may supply the link-level connection to the interface (e.g., your corporate Ethernet LAN), or you may pay a telecommunications company to provide the connection (e.g. dial-up telephone, ISDN, dedicated circuit, etc.), depending on whether the router is located at your site or at the ISP site.

Through the world-wide telephony network the Internet is accessible from almost anywhere in the world (at low bandwidth, using dial-up telephone lines, as well as cellular and satellite telephony). Higher bandwidth and/or dedicated circuit access to the ISP router depends on the capabilities of local telecommunications companies (and soon, almost certainly, cable TV companies). These connections from the ISP router to your site are sometimes called "tail circuits".

Figure 1 illustrates several strategies for connecting to the Internet.

2. A "window system" is the software that allows applications to provide user interaction through "windows" (areas of the screen reserved for a particular application), keyboard, mouse, etc. The window system provides a hardware independent means for applications to manage their own "screen". Older window systems like SunView, MS Windows, Macintosh windows were typically part of the operating system. The X-Window system is a network distributed and platform independent window system that runs on most Unix systems.

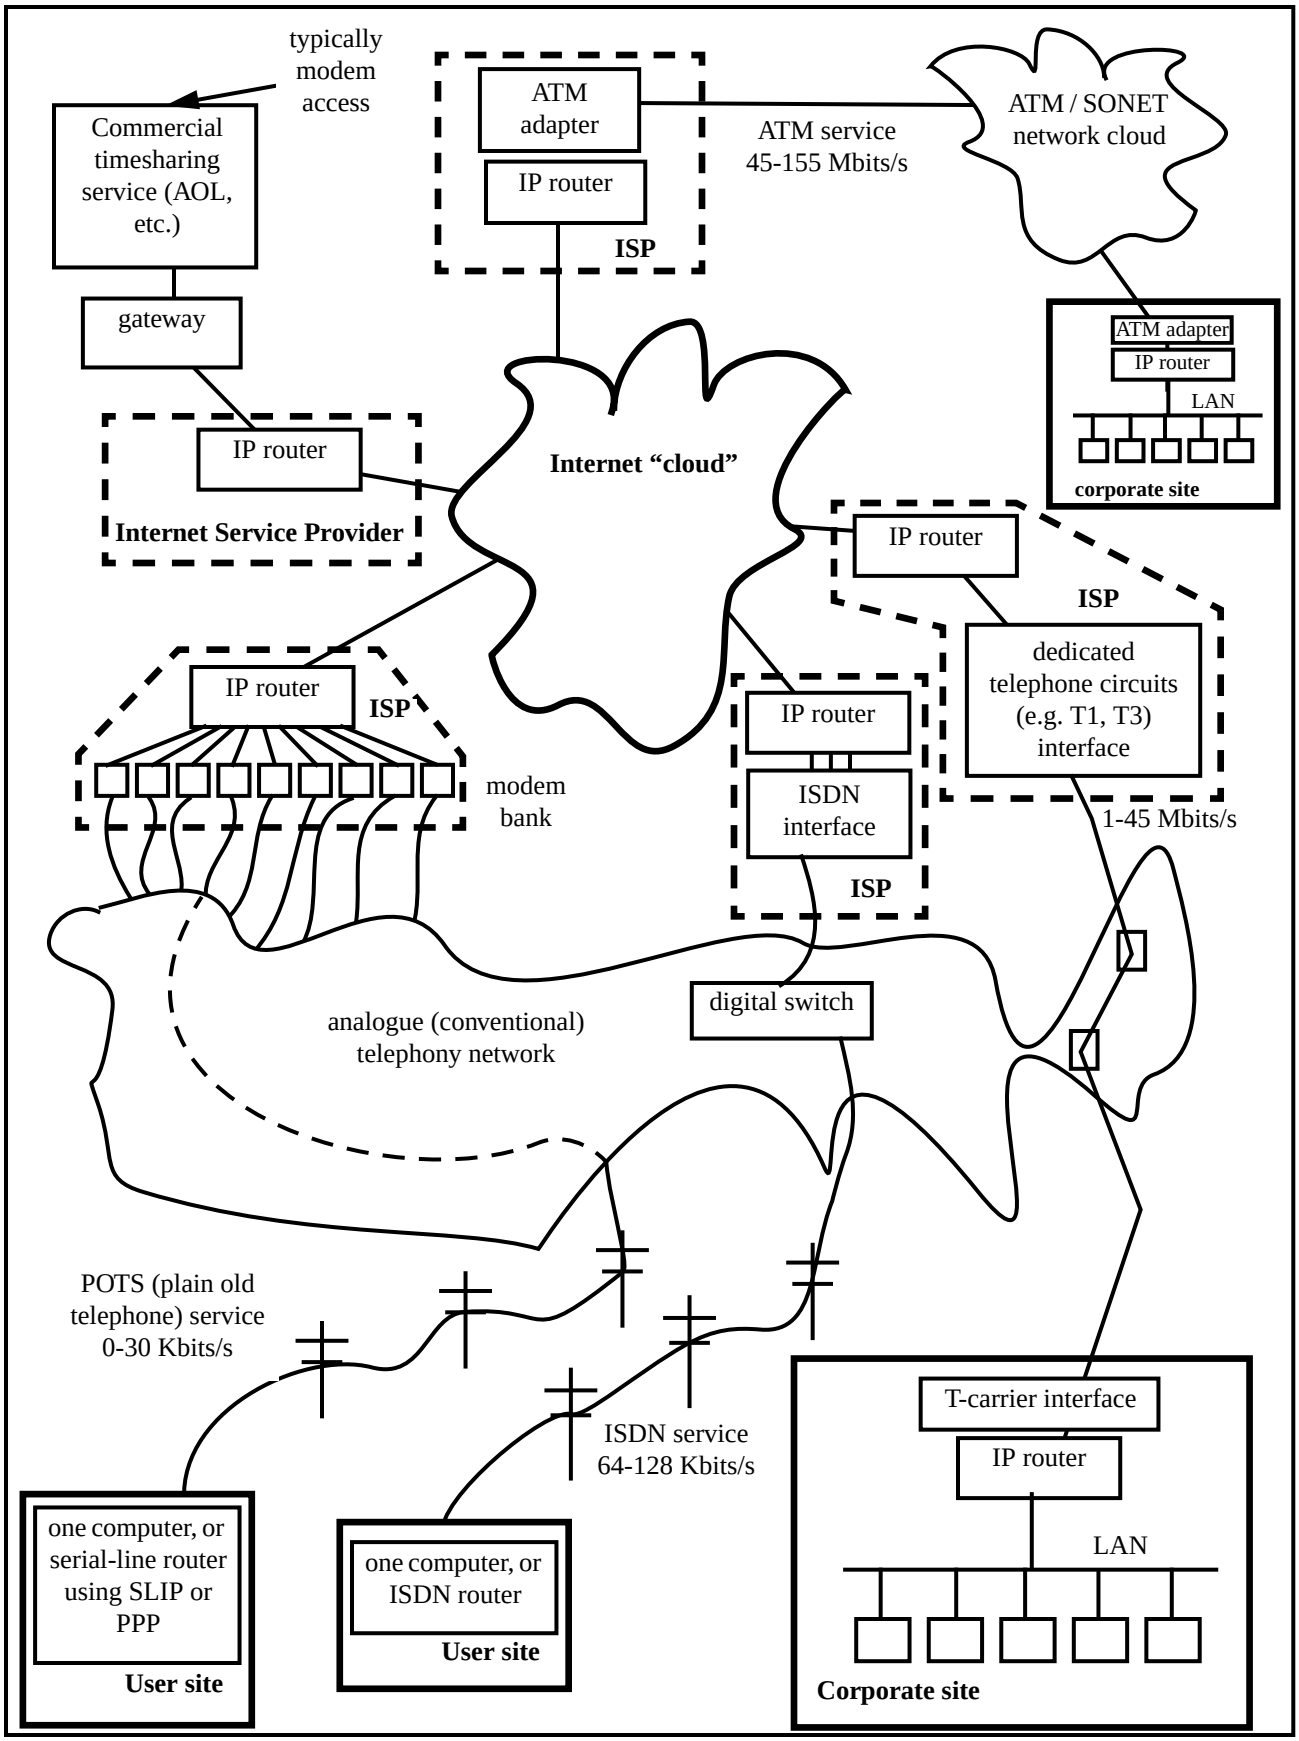

Figure 1 There are many ways that Service Provides can connect users to the Internet.

5.0 Security in the Internet Environment

Security is usually approached from the point of view of identifying the threats (methods of attack) and the associated risks (potential consequences of a successful attack), and then designing protective measures to address these issues (both thwarting attacks, and limiting the damage resulting from a successful attack).

Typical risks include denial of service, unauthorized access to information, unauthorized execution of an action (e.g. transaction), modification of legitimate transactions and data, etc.

While most people tend to think of threats as originating from outside the organization or legitimate community of users, most security breaches (something like 70%) occur from within. However, for the types of information and users associated with TSINs, the main threats are likely to be external, and that is the issue that we focus on here.

Security issues arise, and are dealt with, at several places in the network and attached systems. Very generally speaking, "security" addresses different threats at different "levels" of the overall architecture: the network (and all its many components), the end systems (server and client machines), and the applications themselves. Traditional security objectives include confidentiality, integrity, and availability, as well as legitimate use (of resources). Each of these has associated threats, though the threats differ somewhat depending on the part of the architecture that you are considering. Without undertaking an exhaustive catalogue (see [Ford-95] and [Stallings-95a]), network-level threats, that is threats that might originate within the network, include, for example, denial-of-service, eavesdropping, integrity violation (modification of data), masquerade, replay, and service spoofing. System-level threats include several of the network-level threats plus things like authorization violation, illegitimate use, indiscretions by personnel, theft, trapdoors and Trojan horses (deliberate security flaws in software), etc. Additional application level threats include service spoofing (pretending to be a legitimate service, when you are not), repudiation (of a previous legitimate transaction), etc. We will not address all of these here, but attempt to give a flavor of the kinds of security threat countermeasures that are in use and being developed.

Security in the Internet itself involves the reliable operation of the routers and services essential to the operation of the transport level of the network. These elements are, relatively speaking, few in number and operated by network professionals whose full-time job is to ensure their correct operation. While no human organization or software system is free of bugs that can produce security weaknesses, the single purpose software of routers and the constant attention of people who are aware of the issues, consequences, and solutions, make the Internet infrastructure an infrequently successful target of security-oriented attacks. Successful attacks at this level would probably manifest themselves primarily as service interruptions, which are rare due to any cause, although "wiretapping" and spoofing[3] are also potential consequences of successful attacks on the Internet infrastructure.

3. "Spoofing" is an attack that involves an intruder counterfeiting a valid address in order to get through a firewall that controls access based on addresses.

With respect to specific applications (like the TSIN database servers), most Internet-based security attacks are likely to be indirect. That is, the object of attack is the computer system that runs the application and stores the data, rather than the application itself. The most common mode of this type of attack is through heavily used services (servers), such as mail and ftp, that provide external access to the system.

There are many reasons why indirect attacks are the most common, but it usually comes down to the fact that commonly configured computing systems offer many possible targets, many of which can be replicated and studied by would-be assailants. The targets are typically the servers that provide the common system services of computers attached to the Internet: mail, ftp, remote login, etc. Apart from identifying and correcting all of the design and implementation problems of all of these complex programs (a large and on-going effort by many people in the commercial and academic computer science communities), the best way to reduce vulnerability is to limit the targets or limit the access to them.

The best way to limit targets is to simplify the system. If being a TSIN data server is the only function of a system, then many of the “standard” (and well studied) services do not need to be run, and therefore will not present targets for attack.

The second common way to reduce vulnerability is to limit access by restricting what servers can be accessed from the Internet, or by permitting only known external systems to access the services. These safeguards are accomplished by “firewall” methods that operate most commonly on a router that sits between the end-user system (e.g., the TSIN data server) and the rest of the Internet, but firewall techniques can also be usefully applied on end-user systems themselves. Firewalls are “filters” that screen incoming traffic based on type (what service is being addressed on the inside), originator (large classes of possible originating hosts may be disallowed), receiver, etc. They provide a very effective set of techniques, but do not protect. for example, against all types of address spoofing (a mode of attack where the assailant masquerades as a “friendly” system).

Many security breaches at the system level (both from within and from outside) are due to poor system administration practices. This issue is addressed by an increasingly large body of well publicized knowledge on safe and standard practice for running a system so that it does not inadvertently expose itself to attack, together with operating policy to ensure that safe practice is followed. This kind of information is available from books, articles, professional training, and Internet standards community documents. For example, see [CERT], [Curry-92], [SSH],[Garfinkel-91], [Harvard], [RFC-1244], and [Spafford].

Finally we come to the direct attack threats that most people associate with security issues: the compromise of the data stream and/or the application itself.

Techniques to address threats at this level have developed rapidly over the past few years due to the increasing interest in providing financial services over the Internet. The types of services and techniques that are now available for “end-to-end” security address the issues of authenticated access, data integrity, confidentially, and (potentially) non-repudiation of transactions. The advantage of end-to-end security techniques is that for the aforementioned objectives, one does not have to depend on (trust) security at lower levels. (Which does not mean that lower-level security is not needed. If you network does not function to your data is destroyed on disk, end-to-end security is a moot point.) Most of the end-to-end techniques being adopted by the Internet financial community are based on a technology known as “public key cryptography” and the associated public key certificate infrastructure.

5.1 How do Public Key Certificates Address Internet Security?

Public Key Certificates are a combination of two technologies - public key cryptography and network directory services - that are now being used as the basis of Internet financial transaction security by a wide variety of commercial concerns (e.g. most commercial WWW browsers - Netscape and Spyglass - Mastercard, Visa, Microsoft, etc. - see [SSLNews]).

Public key cryptography uses what is known as asymmetric key-based encryption: Two associated encryption keys are generated as a pair, and documents encrypted with one key are decrypted with the other key. “Conventional” (or “secret key”) cryptography uses the same key to encrypt and decrypt, and is therefore called symmetric encryption. The DES standard for “bulk” encryption is a symmetric scheme. See the [RSA] for an introduction to public key cryptography and the excellent book by Warwick Ford [Ford-95] for an in-depth coverage of the whole field.

The asymmetric nature of public key cryptography (PKC) leads to some interesting properties that are the reason for its value and success in securing wide area network-based transactions. The most common way to use PKC is for the originator of a document to encrypt that document with one of the two keys (usually called the “private” key). The encrypted document and the other key (called the “public” key) are then both distributed widely. The public key is then used to decrypt the document. Successful decryption with the public key means that the document could only have been encrypted with the corresponding private key (and therefore by the holder of that key), and that it must be exactly the same as the document that was originally encrypted.

This technique, then, accomplishes two important goals: It verifies the integrity of the document in-hand (it must be an exact copy of the original) and it verifies the authenticity (only the holder of the private key corresponding to the public key that decrypted the document could have “signed” the original).

In a useful variation of this scheme, documents encrypted using a public key can only be decrypted using the corresponding private key, thus providing a means for anyone in possession of the public key to correspond in guaranteed privacy with the holder of the private key. The principle (described above) and the practice, however, are somewhat different. Public key cryptography is computationally expensive, and so is usually used only for “small” messages. Typical examples of such messages include “digital signatures” (small codes that verify the integrity of a message or document, when used with public key certificates described below, but do not disguise it contents), or the secret key for the bulk en/decryption of a document.

Public key technology is useful in and of itself, but a number of other concepts and supporting infrastructure are needed to make it the useful and powerful tool that it is today, and this brings us to certificates.

The pivotal issue is how public keys are generated, irrefutably associated with a know entity (person, corporation, etc.), and then distributed so that the originator of a document can be both identified and verified.

One important concept is that of Certification Authorities, and how they work. (See [Kent] and [Ford-95].) The problem to be solved is how you associate a public key with an entity (e.g. person), and then distribute that public key so that the relationship (of entity to key) cannot be counterfeited. This is accomplished in the following way. A Certification Authority (“CA”) is a trusted third party that independently verifies the identity of an individual (or other “entity”), issues a public/private key

pair, and then makes publicly available the identity and the corresponding public key of the registrant. The individual, of course, must keep the private key completely confidential.[4]

The identity and the public key are placed in a document that the CA then "signs" with its private key[5]. This document - the "certificate" - is then made publicly available, typically through some sort of network accessible database.

The recipient of a "signed" document retrieves the public key certificate of the presumed author, verifies (through the use of decryption) this certificate with the public key of the CA, extracts the author identity and public key, and then finally verifies the original document. Successful decryption of the original document (or its seal) now verifies its authenticity and the identity of the author. This, then, is the basis of secure, authenticated Internet transactions in an open, scalable, and public network.

There are two other important details to address in order to complete the conceptual framework of public key certificates.

Public key certificates are distributed from network accessible databases. While many variations are possible, the current trend is to use the X.500 network directory server. The primary use of directory servers is to permit e-mail addresses, postal addresses, telephone numbers, etc. to be found for individuals. However, X.500 servers can also supply public key certificates (in a standard form, called an X.509 certificate). These servers are available from anywhere on the Internet and are typically maintained by the parent organization (e.g. a company, university, government laboratory, etc.). This infrastructure, then, provides for the wide distribution of identity-public key information.

The second remaining detail is who issues the certificate and verifies identity. Currently this is done by certification authorities typically associated with a particular community of users. So, for example, Apple Computer has a distributed collaboration environment in which "strong" user authentication is required. To support this activity Apple runs a certification authority. In order to register with this CA (and be issued a public/private key pair) an individual must present documents such as a birth certificate and drivers license to a Notary Public, who signs the registration form. On this basis, the CA provides a certain level of "guarantee" for the identity of the individual associated with the public key. (See [Baum-94] for a detailed discussion of certification authorities and their legal and sociological implications.) RSA, Inc. - the provider of most commercial implementations of public key technology - also runs a CA. Large scale CAs in the future will probably be run by the U. S.

---

4. CAs are frequently associated with a "hierarchical" trust model. That is, trust of the key-entity relationship in the certificate is provided by the guarantee of a trusted third party, who, in turn, may be guaranteed by a "higher" level CA, etc. An alternative is the "web of trust" model of PGP, where each certificate is signed by a collection of "friends", and to establish trust you look for a certificate signer that you know and trust. See [Stallings-95a]. (PGP, or "Pretty Good Privacy", is a widely distributed, public key based system for providing secure e-mail using conventional e-mail systems. See [Stallings-95b].) There is, however, no fundamental difference between the two models: they both end up establishing a "chain of certificates" - the means of tracing an identify claim back to the originator of the claim (the original issuer of the certificate).

5. To "sign" a message means to encrypt the message, or some unique representation of the message, with the private key of the originator. To "seal" a message is to provide an irrefutable integrity check (typically some type of hash code of the message called a "message digest") that is attached to the message as an "appendix". The message digest can easily be recomputed for the received message and compared to the decrypted digest contained in the appendix in order to verify the authenticity of the document. This is a common way to send documents whose value is related to their authenticity (e.g. security bulletins). The digest has the property that no part of the message can be changed with out generating a different digest. Neither necessarily implies that the message itself has been encrypted. (See [Ford-95].)

Postal Service. The USPS CA will provide several levels of certification (identity guarantees) ranging from anonymous to biometrically verified (e.g. retina scan). USPS sees this, and a variety of other cryptographically-based services, as one of its major commercial services in the future. (See [USPS].)

Given the techniques of public key certificates and the infrastructure of network directory servers, security services and protocols can be developed to address the application-level security issues. The services define how an application achieves security, and the protocols specify how to use the basic information and verification techniques (i.e. what information is presented, when, in what order, etc.) in order to accomplish "strong" user (and server) authentication, confidential and verifiable transactions, etc. Description of the services and protocols is beyond the scope of this overview, but examples include the Generic Security Service Application Program Interface (with "Simple Public-Key Mechanism - SPKM) ([GSSAPI]), the Secure Socket Layer protocol (see [SSLProtocol]), Privacy Enhanced Mail ([Kent]), MIME Object Security Services ([MOSS]), and The Secure HyperText Transfer Protocol (and associated services)([SHTTP]).

Two more comments are needed in the interest of completeness. The first is that public key cryptography is almost always used in conjunction with symmetric cryptography (e.g., DES). For the exchange of confidential information PKC is typically used for the authentication and then to provide secure messaging to exchange temporary secret keys for DES encryption/decryption. These temporary keys are sometimes called "session keys" and are used by the two parties for a limited period of time to encrypt and decrypt documents and data. The reason for this is that the DES is much more computationally efficient for "bulk" encryption of data streams.

An important final note is the following. As we have seen, security is a many faceted thing, involving policy, operational practice, network servers, client programs, etc. At the most elementary level there are two types of computer code whose correctness is essential for the underlying concepts to provide security. The first is the cryptographic algorithms themselves (e.g. DES) and the second is the protocols that use these basic algorithms to accomplish conditions like user authentication. Experience has shown that good implementations are those that have stood the test of time (and attacks), and have had their weaknesses identified and corrected. Weaknesses can occur in the protocols (somewhere in a sequence of operations, a critical piece of information is inadvertently exposed) or in the computer code that implements the algorithms and protocols. A good example of this process was reported in a article in the New York Times [NYT-9-19-95]. The Secure Sockets Layer [SSL] used by Netscape for their WWW transaction security is a protocol. Part of the implementation of this protocol involves generating encryption keys on the fly (session keys). Such keys require a "seed" number to generate. If the seed is not a sufficiently random number, it may be possible to break the encryption. The Netscape implementation made the mistake of not using a sufficiently random variable, thus making it possible to break the encryption based on the keys generated with those seeds. Two Computer Science graduate students discovered the weakness, broke the encryption, and announced the fact on an Internet bulletin board. Netscape has now changed the way it generates seeds, carefully re-examined the rest of its code, and a more mature and unbreakable security is the result. In other words, their implementation will have moved toward the maturity that makes for really strong security.

6.0 Cost of Access to the Internet

Internet access service is now a commodity. Thousands of service providers supply service to millions of users. The two principal ways to gain access to the Internet are through direct connection via an Internet Service Provider, and through indirect connection via a timesharing service.

A random sampling of a few dozen Internet service providers from a list of 1300 service providers listed at one WWW site (http://thelist.com/) gave the (rough average) results indicated below (Internet access only, not telecommunications cost).

28.8 Kb/s dial up service	$27/mo
ISDN dial-up service (64 Kb/sec)	$60/mo
56 Kb/sec, dedicated	$300/mo
T1 (1.15 Mb/sec) dedicated	$1300/mo

Recall from the discussion above that this is one component of three required for Internet access: Internet Service Provider, local telecommunications access, and a local piece of hardware to connect from the telecommunication connection to your computers.

The local telecommunications access varies widely from state to state (depending on the state regulatory situation), but in California, for example, a very rough approximation is that the telecommunication costs are about equivalent to the ISP costs.

The hardware required ranges from a few hundred to a few thousand dollars, and within the first year will generally be small compared to the combined service and telecomm costs.

Many users of the Internet gain access to the Internet through VARs that run timesharing systems. These timesharing systems tend to be thematic (oriented toward a particular market segment) and typically supply restricted access to the Internet (via the WWW, telnet, and ftp). Delphi, America Online, and Prodigy are such VARs, and Microsoft and AT&T have recently announced similar services. By way of example [NYT-9-20-95], AT&T's new "Business Network" provides access to business bulletin boards and "access to the World Wide Web via Netscape Communications Navigator software that is included in the package, Internet chat groups and discussion groups related to business topics orchestrated by AT&T." "The service will cost $39.95 a month for the first 10 hours of use and $2.95 for each additional hour." The Microsoft offering is similarly structured. These services can usually be accessed through the Internet, but the most common way is by dial-up access to the timesharing system (and from there through a gateway to the Internet). AT&T is also an Internet Service Provider through its "AT&T World Net Service".

7.0 Reliability and Availability of Internet Service

For large sites (universities and corporations) that deal directly with "regional" Internet Service Providers, the Internet is highly reliable. Part of the reason for the reliability is that it is not difficult to configure alternate paths to and through the Internet "cloud". (Multiple paths through the cloud are an automatic part of the Internet operation, and reflect the fact that the Internet technology was originally designed as a "survivable" network for military use.) Alternate paths can be provided into the Internet cloud, and even (for a "small" number of sites) around the cloud (i.e. a "private" internet). These possibilities are illustrated in Figure 2.

Availability of Internet bandwidth and services are two different, and frequently confused issues. Various pieces of the Internet do get congested, however for most application the transport protocols cause this to be manifest as increasing delays for individual messages. (This is most apparent for interactive terminal sessions, where local key strokes are sent as messages across the INternet, and delay is very obvious.) People mistakenly blame un-available Web service on Internet congestion. This is almost never the case. Web client-server connect failures are usually due to limited resources

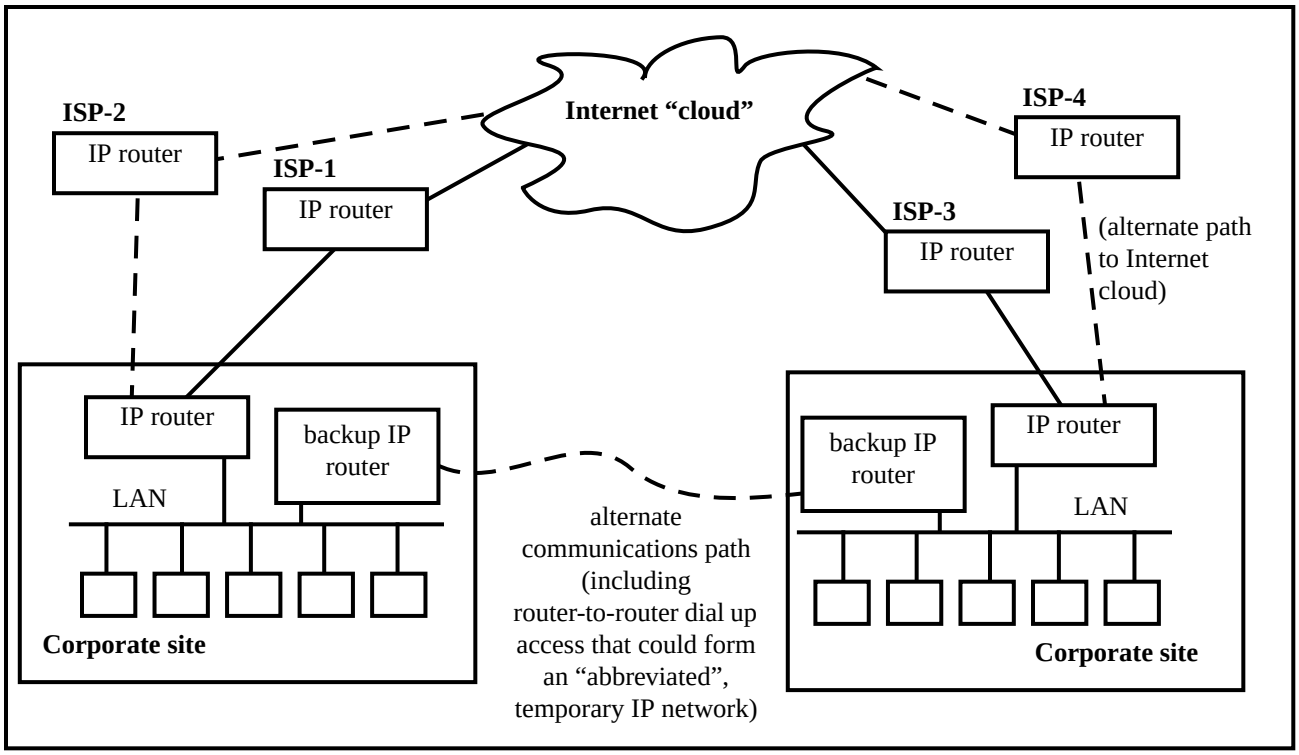

Figure 2 Alternate IP access to sites can be provided for added reliability.

available on the Web server. That is, almost all Web servers are have resource limitations that restrict the number of incoming connections, and it is this contention that causes clients to fail to be able to contact a server.

8.0 Use of the WWW to Provide Access to Transmission Open Access Data[6]

In this section we present a strawman design for both a system architecture and a data management architecture for the TSIN. The purpose of this strawman is principally to provide an "existence proof": That is, to demonstrate at least one approach to providing TSIN service and Transmission Open Access (TOA) data using the WWW and the Internet.

The "requirements" for the design are taken from the discussions at the Sept., 1995 Real-Time Information Networks (RINs) Implementation workshop ([EPRI-95b]), together with some other constraints imposed to simplify the architecture.

The "requirements" from [EPRI-95b] reflected in the design include:

- The TSIN is a reservation / posting system
- Transaction auditing should be provided
- "Trading Partner agreements" will be needed to cover:
  - access to information
  - ability to make reservations (and for what and how much)
- Secure communication of data and transactions should be provided
- "Easy" identification of changed information should be possible
- Each Provider shall be responsible for posting and updating TOA information on TSIN nodes
- Secondary Providers must post TOA information on Primary Provider's TSIN node [not addressed in the strawman design]

How the requirements are met, plus some "implied, derived, or added" conditions to be able to produce a simple, illustrative design are as follows:

- Presentation of the TOA data in will be in several formats:
  - A "human" readable format will accommodate interactive, on-line browsing and interactive transactions through the WWW interface.
  - A "program" readable format will accommodate rapid acquisition and analysis of the TOA data by automated user programs.
  - A "spreadsheet" readable format is intended to accommodate smaller users who don't automatically fetch and analyze data.
- TSIN dial-up access will be provided by Internet Service Providers, commercial timesharing services, or bulletin boards maintained by third party VARs who pick up data from the TSIN nodes and maintain it on their bulletin boards. Dial-up access to TSIN node "bulletin boards" is specifically excluded due to the added security risks of allowing a user "presence" on the node.
- "Up-to-date" information and the ability to detect changes in the TOA data are provided through an incremental update strategy:
  - A "base" file (complete TOA information for a transmission service provider) will be updated at low frequency (e.g., once or twice a day) from the TSP's private databases.

6. Nomenclature (from [EPRI-95b]): "The Transmission System Information Network (TSIN) shall consist of multiple interconnected computer nodes which contain Transmission Open Access (TOA) information from Transmission Providers (Providers), and which permit Transmission Customers (Customers) to access this information."

  - Incremental change files are supplied at high frequency (e.g., once every 5-10 minutes). When applied to the base file in sequence, these incremental change files produce a current view of available transmission capacity.
  - The TSIN node update manager could maintain one completely up-to-date TOA file for "browsing", and for those who only infrequently access the TOS information.
  - Historical information could be kept by archiving the current (updated to that time) TOA file at (say) midnight each day.

- Reservation transactions will be made via the WWW "forms" interface:
  - The ability to make reservations, how much can be reserved, etc., is under the control of a "trading partners" information base. The mechanism of generation and maintenance of this information is not specified, but on the TSIN node it is a carefully protected list of authorized users and their capabilities.
  - WWW forms are the only reservation mechanism that will be supported in order that secure and audit-capable commercial WWW servers intended to support financial transactions may be used by the TSPs for TSIN nodes.
  - Third-party software vendors can supply other client-side (TSIN user) capabilities, such as programs that convert spreadsheets to WWW forms, and interfaces to other user "trading" programs. (In other words, this third party software would generate WWW forms directly from user data and pass these forms to the TSIN WWW server.)

- The data management architecture suggests that the TSIN WWW server will interact with an "update" server that actually manages the TOA data from the TSP. The update server handles loading TSP information onto the TSIN node, manages the base and incremental TOA data files, automatically generates the html documents that describe and index the TOA data files, handles reservation transactions, enforces trading partner conditions, provides audit logs, etc. This functionality is "common" to trading systems in the financial industry. For example see the Lombard demonstrations for their Internet-based, on-line securities trading system. [Lombard]

- An ftp and/or gopher server can easily be configured to provide file-oriented access to the TOA data. ftp servers are among the most mature Internet servers, and have evolved to a relatively high level of security.

- Non-Internet access to the TOA data could be provided through a mail "list" server. (A mail server that is designed to respond automatically to requests for files.) This server is not shown in the architecture diagram.

The strawman TSIN node "security and data flow architecture" is illustrated in Figure 3, and the data management architecture is illustrated in Figure 4.

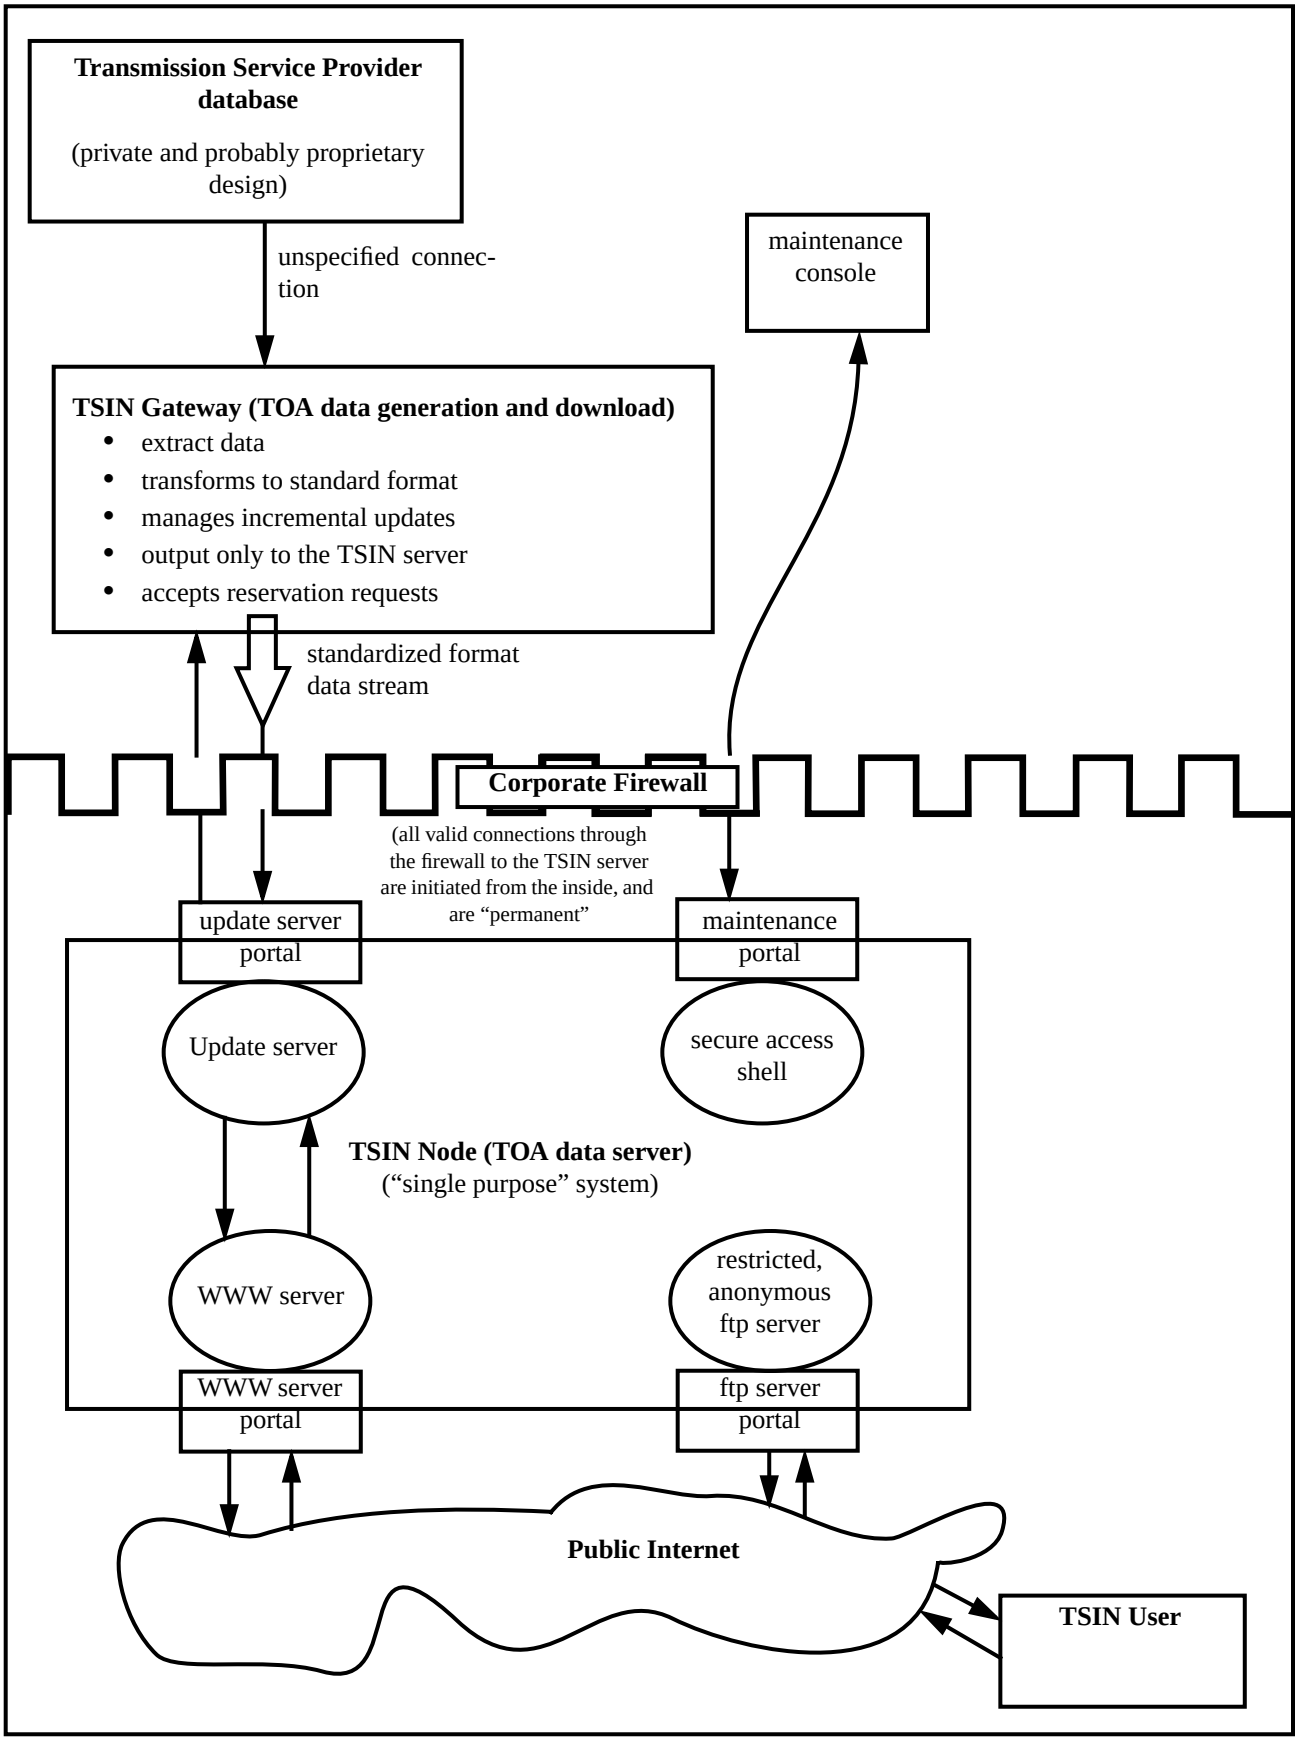

Figure 3 Strawman TSIN system security and data flow architecture.

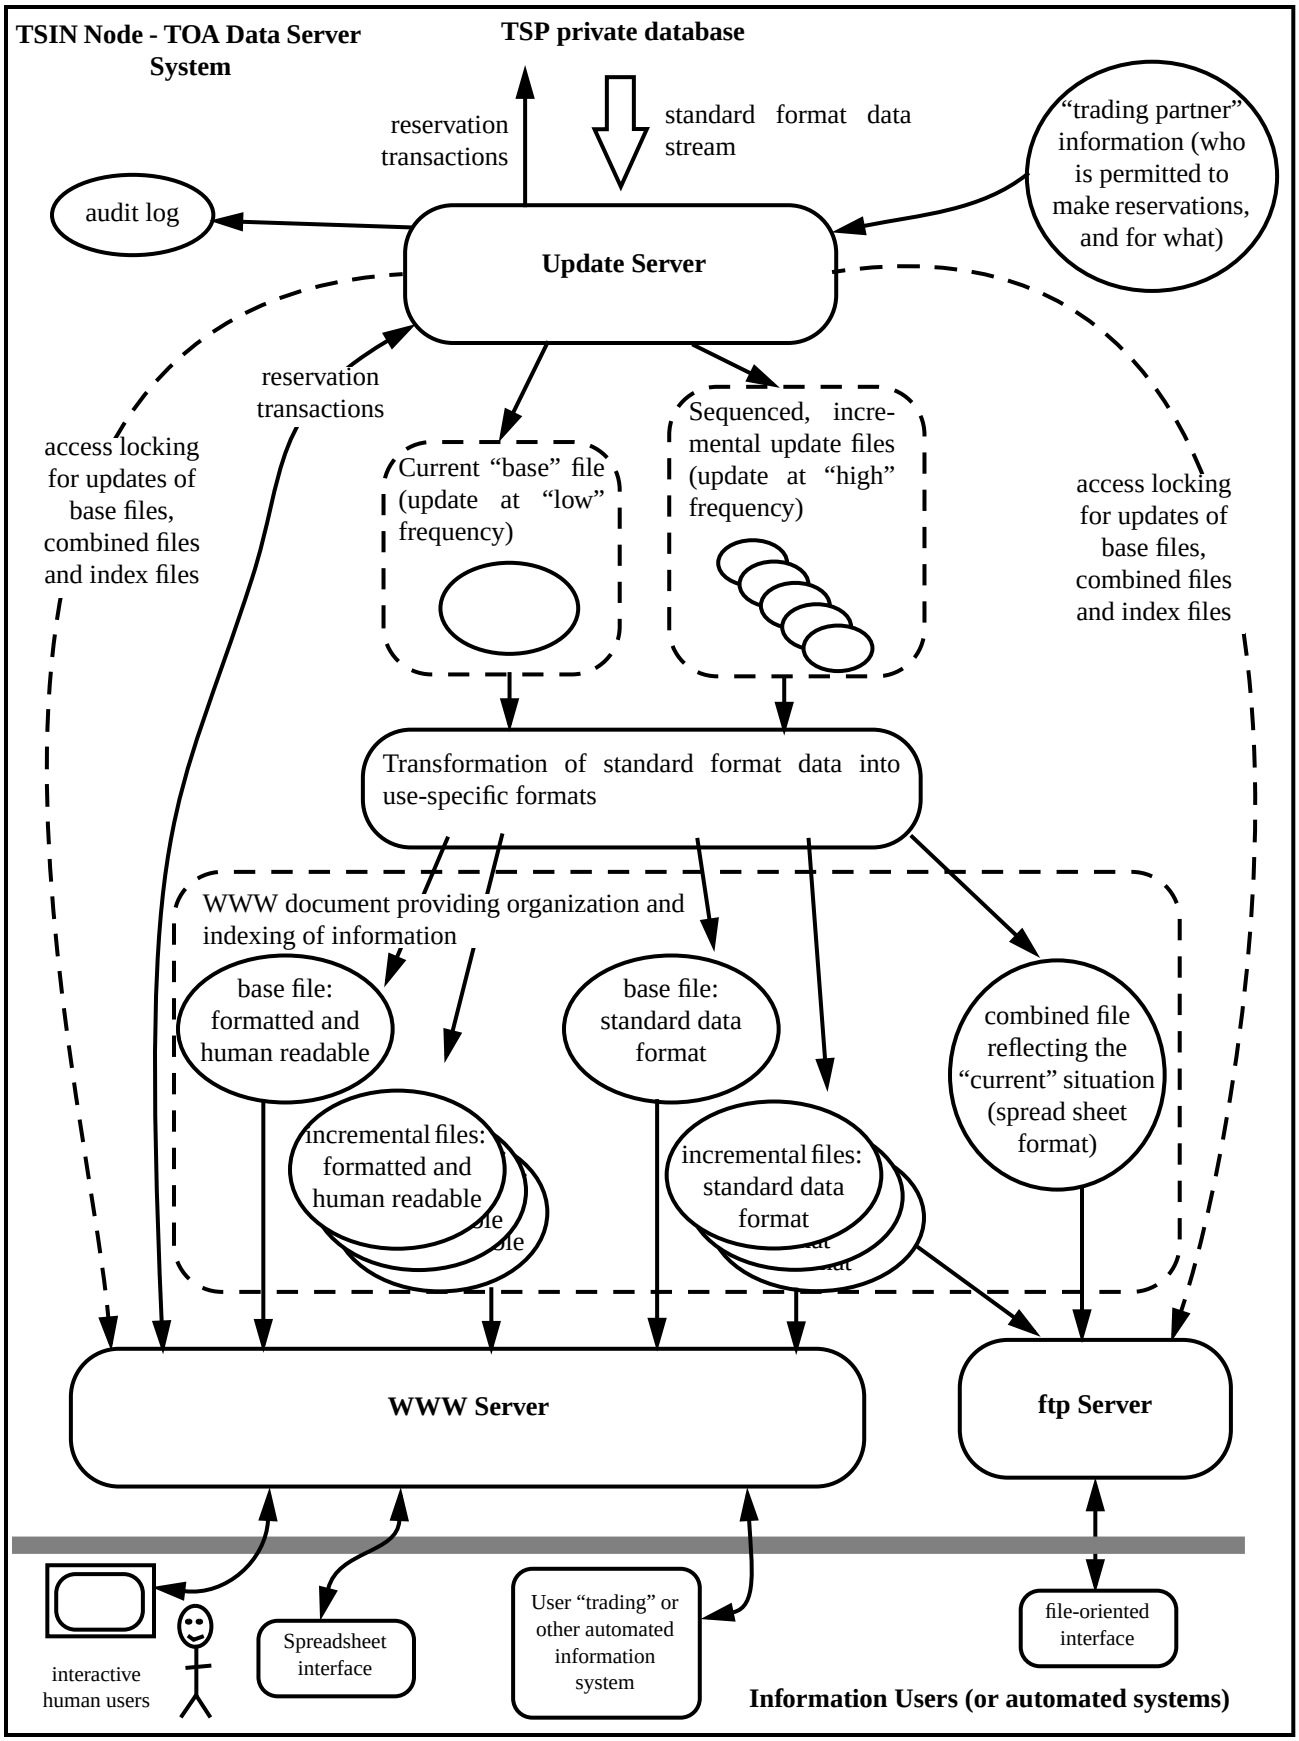

Figure 4 **Strawman TSIN data management architecture.**

9.0 Electronic Data Interchange (EDI / X.12) and the Internet

This section is included mainly to illustrate some of the progress in the migration of one well-established electronic commerce mechanism toward the use of the Internet. There is no particular intention to imply (or not) any connection between EDI and the TSIN environment.

Electronic Data Interchange (EDI) is one element an architecture for conducting electronic commerce (EC). These EC architectures are most commonly focused on addressing complete business processes such as the life cycle of material acquisition. To this end EDI provides or facilitates: RFP distribution, response evaluation, P. O. generation (vendor orders), material tracking, and electronic payment. A typical process model and implementation are given in a report describing a strategy for the Federal Government to use EDI and the Internet for all procurement. (See [ECAT].) The ECAT process model is illustrated in Figure 5.

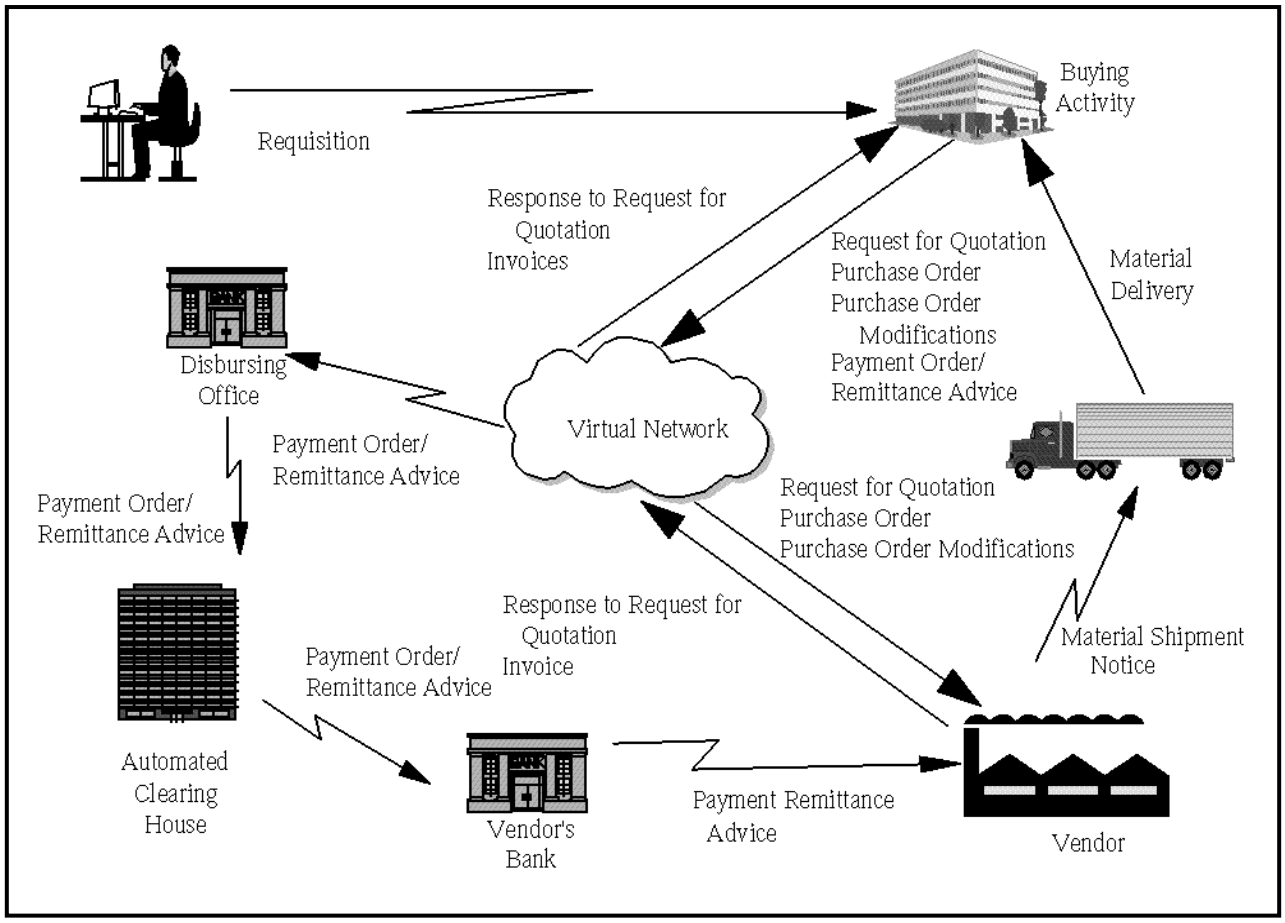

Figure 5 Procurement process model used to define the ECAT architecture. (From [ECAT])

The ECAT implements this process model through an architecture that involves the concepts of value added networks (VANS) and virtual networks (VN). A VN is a collection of heterogeneous network technologies that are integrated through network entry points (NEP) that allow messaging between all of the originators (e.g. government agencies) and the VANs. The VANs are established commercial concerns that provide a variety of services including EDI message forwarding and format conversion.

In the traditional EDI scenario, the VANs provide access via proprietary networks and services that interact with a restricted set of TPs. The ECAT architecture is illustrated in Figure 6, and is an attempt

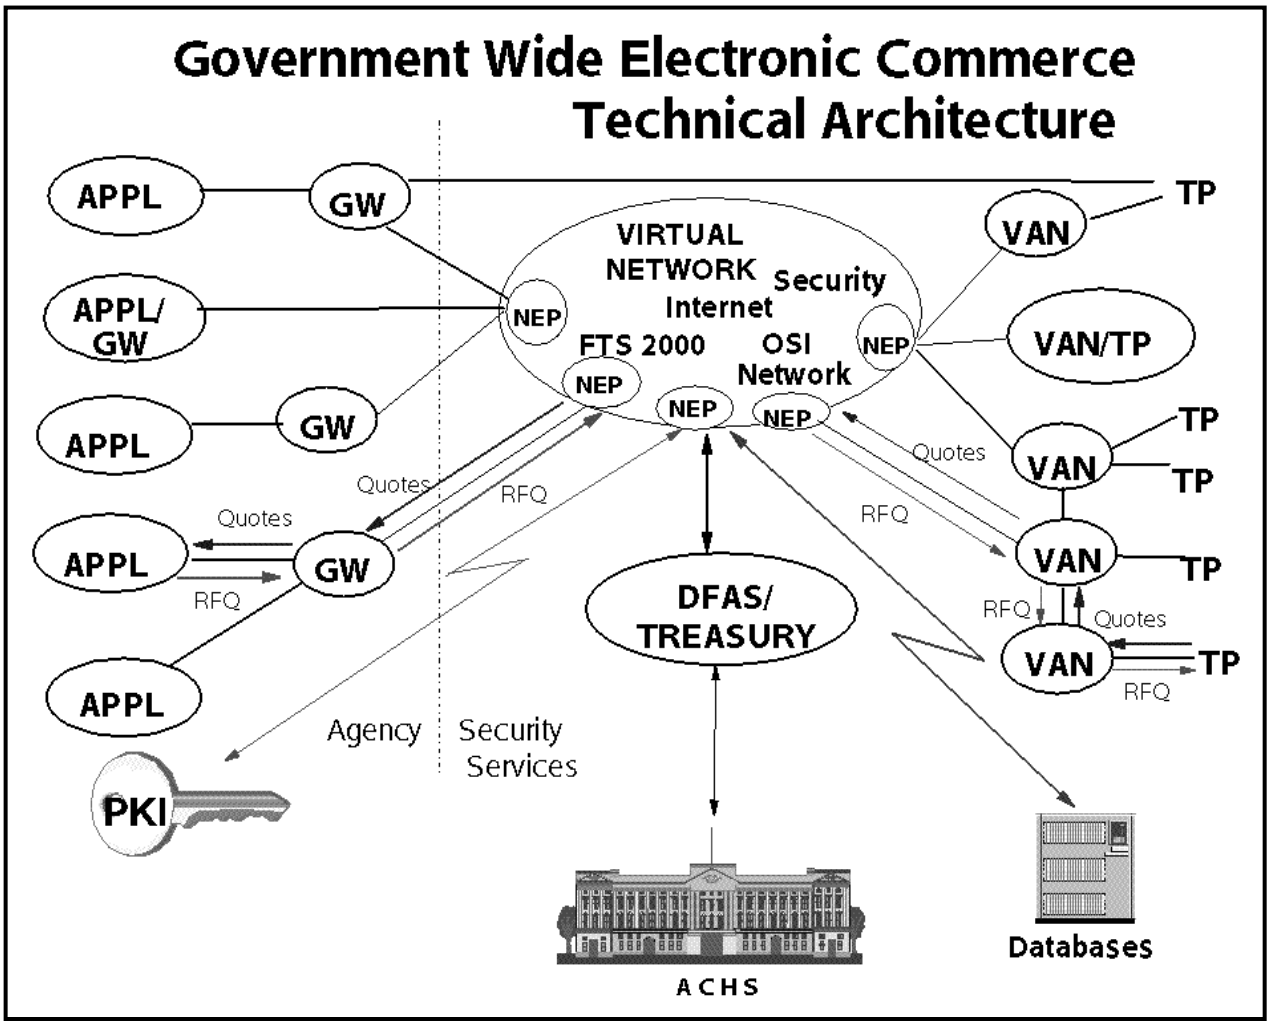

Figure 6 **ECAT architecture.**

In the figure: APPL= User/agency applications GW Gateway; NEP = Network entry point; ACHS = Automated clearinghouse system; VAN = Value-added network; TP = Trading partner (vendor); OSI = E-mail, X.400, X.435, X.500, FTAM; Internet = FTP, SMTP; PKI = Public key infrastructure. (From [ECAT])

to bring some uniformity and openness to this environment. The architecture provides users with the ability to use the Internet as a single point of access to all VANs, and provides application-level gateways to homogenize the EDI message transport environment. The use of the Internet in the ECAT architecture is intended to provide “universal” access to the VANs (rather than individually via a collection of single-purpose communication links), relieving the purchaser of the need to “call” many different VANs in order to conduct the full spectrum of business. Operationally the NEPs might forward e-mail-based EDI messages to the VANs, who then may translate it to the format needed by the specific business / trading partner. A possible end-game of the ECAT approach is that businesses will provide their own Internet access, accepting EDI messages directly. The role of the VANs may well transform to more that of a broker which, for example, collections and provides information on all businesses dealing with a certain product line (e.g. office supplies) so that the buyer does not have

to identify and query many different sites to examine catalogues, get price quotes prior to placing orders, etc.

While the TSIN process model is different from the more traditional procurement model, there are some parallels, and several elements and issues of the corresponding architectures are similar. There are also significant differences that arise from the “real-time” nature of he transmission capacity transactions compared with the more “leisurely” pace of material acquisition systems that use EDI.

However, parallels exist, both in the nature of the process and the issues of bringing business operations on-line to the Internet. (“EDI Meets the Internet: Frequently Asked Questions about Electronic Data Interchange (EDI) on the Internet.”([EDI-INet]) is a document that describes the Internet to the EDI community, and presents some of the Internet-related issues in a fashion not unlike this paper.)

10.0 Acknowledgments

I express appreciation for the following valuable assistance: Mary Anne Scott, Office of Energy Research, Office of Computation and Technology Research, Mathematical, Information, and Computational Sciences (MICS) Division, of the U. S. Department of Energy provided the motivation for this white paper. Robert Fink, David Stevens, and Frank Olken of LBNL provided comments that increased its clarity. Warwick Ford of Bell Northern Research provided several useful comments on the security sections. Any omissions or errors are, of course, my responsibility - WEJ.

11.0 References and Notes

Baum-94 M. Baum, “Federal Certification Authority Liability and Policy: Law and Policy of Certificate-Based Public Key and Digital Signatures”, Published by: U. S. Dept. of Commerce, NIST, June, 1994.

CERT From ftp://ftp.cert.org/pub/cert_faq:

> “The CERT Coordination Center is the organization that grew from the computer emergency response team formed by the Defense Advanced Research Projects Agency (DARPA) in November 1988 in response to the needs exhibited during the Internet worm incident. The CERT charter is to work with the Internet community to facilitate its response to computer security events involving Internet hosts, to take proactive steps to raise the community’s awareness of computer security issues, and to conduct research targeted at improving the security of existing systems.”

Comer Douglas Comer, *Internetworking with TCP/IP,* Volumes 1, 2, and 3, Prentice Hall

These books are the “standard” textbooks on TCP/IP. See, for example http://www.prenhall.com/~rich/013/216986/21698-6.html

Curry-92 David Curry, *UNIX System Security: A Guide for Users and System Administrators*. Reading, MA: Addison-Wesley Publishing Co., Inc., 1992. (ISBN 0-201-56327-4)

ECAT Federal Electronic Commerce Acquisition Team, *Streamlining Procurement Through Electronic Commerce*, available from http://snad.ncsl.nist.gov/dartg/edi/arch.html. October 13, 1994, Federal Electronic Commerce Acquisition Team, Skyline 4, Suite 400 5113 Leesburg Pike Falls Church, Virginia 22041.Tel: (703) 681-0369, FAX: (703) 681-0362.

This report is a good introduction to the use of EDI within a complete electronic commerce architecture. It describes a strategy for converting all government procurement to use EDI.

> “The recommended architecture and underlying rationale consists of the following fundamental components:
>
> - A single means of supplier registration to do business electronically with the Federal government including a standardized trading partner agreement embodying the “rules of the road”
> - A standard method of implementing the electronic data interchange (EDI) transaction formats used in the United States [currently those approved by the American National Standards Institute (ANSI) Accredited Standards Committee(ASC) X12]

- Existing agency-managed procurement systems modified to generate standard EDI ASC X12 transactions (i.e., agencies would modify their existing systems to feed data in "flat file format" to a commercial off-the-shelf software package called a translator that generates the ASC X12 transaction)
- A "virtual network" connecting agency standardized transactions to facilities where value-added networks (VANs) or other entities can access them
- A standard agreement between the government and the VANs that support the government and its trading partners
- A standards-based system that gives agency procurement staff access to government data bases supporting their operations
- The use of electronic funds transfer (EFT) as the principal method of payment and the development of a supportive EFT architecture."

EDI-INet W. Houser, J. Griffin and C. Hage, "EDI Meets the Internet: Frequently Asked Questions about Electronic Data Interchange (EDI) on the Internet." Internet RFC-1865. (Available at ftp://ds.internic.net/rfc/rfc1865.txt)

EPRI-95a Electric Power Research Institute, Power Delivery Group, "Real-Time Information Networks (RINs) Implementation Information". A WWW page at A WWW page at http://www.epri.com/org/pdg/ssos/rin/rininfo.html

EPRI-95b Electric Power Research Institute, Power Delivery Group, "EPRI RIN Working Group Summary Presentations of 'How' Transmission Service Information Networks (TSINs) will Operate". A WWW page at http://www.epri.com/org/pdg/ssos/rin/wrkgrp.html

Ford-95 Warwick Ford, *Computer Communications Security: Principles, Standards, Protocols, and Techniques*, Prentice Hall, Englewood Cliffs, New Jersey, 07632, 1995. ISBN 0-13-799453-2.

Frog David Robertson, William Johnston, and Wing Nip, "Virtual Frog Dissection: Interactive 3D Graphics Via the WWW," Proceedings, The Second International WWW Conference '94: Mosaic and the Web, Chicago, IL (1994). (Available at http://www-itg.lbl.gov/vfrog/WWW.94.paper.html.)

From the paper:

"We have developed a set of techniques for providing interactive 3D graphics via the World Wide Web (WWW) as part of the 'Whole Frog' project. We had three goals: (1) to provide K-12 biology students with the ability to explore the anatomy of a frog with a virtual dissection tool; (2) to show the feasibility of interactive visualization over the Web; and (3) to show the possibility for the Web and its associated browsers to be an easily used and powerful front end for high-performance computing resources."

"We have developed techniques to utilize the Common Gateway Interface (CGI) capability of WWW servers to provide an interactive 3D visualization front end through Web clients. These techniques have been used to make a 'Virtual Frog Dissection Kit'. A student using this kit has the ability to view various parts of a frog from many different angles, and with the different anatomical structures visible or invisible. For example, the student can press 'form' buttons that indicate that he or she wants to view the frog from above, with the exterior and skeleton removed. An advantage to this technique, as opposed to dissecting a real frog, is that undissection is as easy as dissection."

"The kit has a forms -based interface. Form submission results in a call to a CGI script, which in turn contacts a continuously running process on a more powerful machine to accomplish the graphics rendering of a large 3D data set representing the frog and its internal organs. The resulting image is converted to Graphics Interchange Format (GIF) encoding. When that process completes generation of the image, it passes the location of the image file and control back to the script which rewrites the image on the client. While this might sound awkward, the overall process is quite similar to how [conventional] rendering systems work, [where] the image [is] being written into a local frame buffer, or sent across the network as an X-window image."

Also see http://george.lbl.gov/frog .

Garfinkel-91 Simson Garfinkel and Gene Spafford, *Practical UNIX Security.* Sebastopol, CA: O'Reilly & Associates, Inc., 1991. (ISBN 0-937175-72-2)

GSSAPI J. Linn, "Generic Security Service Application Program Interface, Version 2", an Internet Engineering Task Force draft from the Common Authentication Technology Working Group (http://www.ietf.cnri.reston.va.us/html.charters/cat-charter.html). The GSSAPI document is at ftp://ds.internic.net/internet-drafts/draft-ietf-cat-gssv2-03.txt .

From the Abstract:

"The Generic Security Service Application Program Interface (GSS-API), as defined in RFC-1508, provides security services to callers in a generic fashion, supportable with a range of underlying mechanisms and technologies and hence allowing source-level portability of applications to different environments. This specification defines GSS-API services and primitives at a level independent of underlying mechanism and programming language environment, and is to be complemented by other, related specifications:

- documents defining specific parameter bindings for particular language environments
- documents defining token formats, protocols, and procedures to be implemented in order to realize GSS-API services atop particular security mechanisms

This Internet-Draft revises RFC-1508, making specific, incremental changes in response to implementation experience and liaison requests. It is intended, therefore, that this draft or a successor version thereto will become the basis for subsequent progression of the GSS-API specification on the standards track."

Strictly speaking the GSS-API defines a service rather than a protocol. An example of a specific protocol implementing the GSS-API is to be found in "The Kerberos Version 5 GSS-API Mechanism", J. Linn, ftp://ds.internic.net/internet-drafts/draft-ietf-cat-kerb5gss-02.txt .

Harvard Harvard University, "Information Security Handbook", gopher://gopher.harvard.edu:70/00/.vine/providers/oit/Computer_Security_Handbook/Information_Security_Handbook/

From the Background section:

"An Information Security Working Group has been organized to review issues of safekeeping and confidentiality of information resources, identify risks, raise consciousness in the community and, where appropriate, develop policy statements, advisories, and guidelines. The working group has representatives from almost all the schools and major central administration departments. The intention is to build consensus among these groups, promote common definitions, compile good practices and check lists in the form of an Information Security Handbook which will be published and updated as the need arises.

While the effort was initially intended to look at administrative computer systems and the electronic distribution of data to the desktop, it was felt that the security issues of paper files, library, and research data could not be excluded. Many of the security practices recommended in the handbook are already standard practice for paper documents; they need to be extended to electronic forms of information as well. Moreover, the integration of systems across mainframes, minicomputers, microcomputers and networks makes it impossible to separate many of the concerns by application type. Security issues must be considered across many environments and media, including paper, which are increasingly shared among a heterogeneous community of users. Many of the people involved in this working group and in the University at large have cross functional responsibility and must look at security issues across their entire organizations."

Hedrick Charles L. Hedrick "Introduction to the Internet Protocols" (Available from http://www.aetc.af.mil/tutorials/ipintro.html or ftp://nic.merit.edu/introducing.the.internet/intro.to.ip)

This document is a brief introduction to the Internet networking protocols (TCP/IP). It includes a summary of the facilities available and brief descriptions of the major protocols in the family.

IRC IRC stands for "Internet Relay Chat".

From http://www.main.com/dms/irc.html :

"It was originally written by Jarkko Oikarinen (jto@tolsun.oulu.fi) in 1988. Since starting in Finland, it has been used in over 60 countries around the world. It was designed as a replacement for the "talk" program but has become much more than that. IRC is a multi-user chat system, where people convene on "channels" (a virtual place, usually with a topic of conversation) to talk in groups, or privately. IRC is constantly evolving, so the way things to work one week may not be the way they work the next. Read the MOTD (message of the day) every time you use IRC to keep up on any new happenings or server updates. IRC gained international fame during the 1991 Persian Gulf War, where updates from around the world came across the wire, and most irc users who were on-line at the time gathered on a single channel to hear these reports. IRC had similar uses during the coup against Boris Yeltsin in September 1993, where IRC users from Moscow were giving live reports about the unstable situation there."

"IRC works when the user runs a "client" program (usually called 'irc') which connects to the irc network via another program called a "server". Servers exist to pass messages from user to user over the irc network."

Java "Java: A simple, object-oriented, distributed, interpreted, robust, secure, architecture neutral, portable, high-performance, multithreaded, and dynamic language" that provides for interactive applications in the context of the World Wide Web. http://java.sun.com/1.0alpha3/doc/overview/java/index.html More generally, see: http://java.sun.com/

Johnston-95 "Realtime Information Networks for Open Access to Electric Power Transmission Facilities" an informational WWW Page at http://www-itg.lbl.gov/~johnston/EDM/FERC-NOPRA.html

Kehoe Brendan P. Kehoe "Zen and the Art of the Internet: A Beginner's Guide to the Internet" Available from gopher://nic.merit.edu:7043/0/introducing.the.internet/zen.txt or http://sundance.cso.uiuc.edu/Publications/Other/Zen/zen-1.0_toc.html .

This is a bit dated (no information on the WWW) but a classic introduction to the Internet.

Kent Steve Kent, “Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management”, RFC-1422, http://ds.internic.net/rfc/rfc1422.txt .

RFC-1422 establishes the conceptual framework for certification authorities. The hierarchical authority chain defined in the RFC will probably be modified in the next version of the document, but most of the concepts are still valid. From the RFC:

> “This document defines a supporting key management architecture and infrastructure, based on public-key certificate techniques, to provide keying information to message originators and recipients.....
>
> The key management architecture described in this document is compatible with the authentication framework described in CCITT 1988 X.509. This document goes beyond X.509 by establishing procedures and conventions for a key management infrastructure for use with Privacy Enhanced Mail (PEM) and with other protocols, from both the TCP/IP and OSI suites, in the future. There are several motivations for establishing these procedures and conventions (as opposed to relying only on the very general framework outlined in X.509):
>
> - It is important that a certificate management infrastructure for use in the Internet community accommodate a range of clearly-articulated certification policies for both users and organizations in a well-architected fashion. Mechanisms must be provided to enable each user to be aware of the policies governing any certificate which the user may encounter. This requires the introduction and standardization of procedures and conventions that are outside the scope of X.509.
> - The procedures for authenticating originators and recipient in the course of message submission and delivery should be simple, automated and uniform despite the existence of differing certificate management policies. For example, users should not have to engage in careful examination of a complex set of certification relationships in order to evaluate the credibility of a claimed identity.
> - The authentication framework defined by X.509 is designed to operate in the X.500 directory server environment. However X.500 directory servers are not expected to be ubiquitous in the Internet in the near future, so some conventions are adopted to facilitate operation of the key management infrastructure in the near term.
> - Public key cryptosystems are central to the authentication technology of X.509 and those which enjoy the most widespread use are patented in the U.S. Although this certification management scheme is compatible with the use of different digital signature algorithms, it is anticipated that the RSA cryptosystem will be used as the primary signature algorithm in establishing the Internet certification hierarchy. Special license arrangements have been made to facilitate the use of this algorithm in the U.S. portion of Internet environment.”

Lemay-95 Laura Lemay, “Teach Yourself Web Publishing in a Week,” Published by Sams Publishing, 1995. ISBN 0-672-30667-0 (Tel: 1-800-428-5331, http://www.mcp.com/cgi-bin/do-bookstore.cgi)

> “This book covers HTML 2.0, images, sound and video, servers, CGI, forms, and imagemaps. In addition, unlike most other books about HTML, this book will teach you about how to create well-designed, easy to navigate and maintainable Web presentations with information about design, organization, and effective linking.”

Liu-95 Cricket Liu, Jerry Peek, Russ Jones, Bryan Buus & Adrian Nye, *Managing Internet Information Services: World Wide Web, Gopher, FTP, and more*, O'Reilly & Associates, 1995. (http://www.ora.com, Tel: 1-800-889-8969), ISBN: 1-56592-062-7.

"This comprehensive guide describes how to set up information services and make them available over the Internet. It discusses why a company would want to offer Internet services, provides complete coverage of all popular services, and tells how to select which ones to provide. Most of the book describes how to set up Gopher, World Wide Web, FTP, and WAIS servers and email services."

Lombard Lombard Institutional Brokerage, Inc. "Real-Time Trading and Research Information" http://www.lombard.com/

From the introduction:

"Welcome to the Lombard Institutional Brokerage Real-Time Trading and Research Information Center. Our philosophy is simple: 'Through the use of cutting edge technology, we are dedicated to providing our customers in the Internet community with a wide variety of investment options, enhanced investment tools and an unparalleled commitment to customer service...' "

For an example of a trading system interface, specifically see the Lombard demonstration at http://www.lombard.com/Demo/

Mbone Michael R. Macedonia and Donald P. Brutzman, Naval Postgraduate School, "MBone Provides Audio and Video Across the Internet," IEEE COMPUTER magazine, pp. 30-36, April 1994.

From ftp://taurus.cs.nps.navy.mil/pub/mbmg/mbone.html:

"Short for Multicast Backbone, MBone is a virtual network that has been in existence since early 1992. It was named by Steve Casner of the University of Southern California, Information Sciences Institute and originated from an effort to multicast audio and video from meetings of the Internet Engineering Task Force. Today, hundreds of researchers use MBone to develop protocols and applications for group communication. Multicast provides one-to-many and many-to-many network delivery services for applications such as video conferencing and audio where several hosts need to communicate simultaneously. The magic of MBone is that teleconferencing can be done in the hostile world of the Internet where variable packet delivery delays and limited bandwidth play havoc with applications that require some real-time guarantees. Limited experiments demonstrated the feasibility of audio over the ARPAnet as early as 1973. However, only a few years ago, transmitting video across the Internet was considered impossible. Development of effective multicast protocols disproved that widespread opinion. In this respect, MBone is like the proverbial talking dog: It's not so much what the dog has to say that is amazing, it's more that the dog can talk at all!"

"The key network concepts that make MBone possible are IP multicast and real-time stream delivery via adaptive receivers. For example, in addition to the multicast protocols, many MBone applications are using the draft Real-Time Protocol on top of the User Datagram Protocol and Internet Protocol. RTP, being developed by the Audio-Video Transport Working Group of the Internet Engineering Task Force, provides timing and sequencing services, permitting the application to adapt and smooth out network-induced latencies and errors."

Also see the “Mbone Homepage” http://www.best.com/~prince/techinfo/mbone.html and http://www.rpi.edu/Internet/Guides/decemj/itools/cmc-mass-mbone.html

MOSS S. Crocker, N. Freed, J. Galvin, S. Murphy, “MIME Object Security Services.” Internet RFC-1848. (Available at ftp://ds.internic.net/rfc/rfc1848.txt)

“This document defines MIME Object Security Services (MOSS), a protocol that uses the multipart/signed and multipart/encrypted framework to apply digital signature and encryption services to MIME objects. The services are offered through the use of end-to-end cryptography between an originator and a recipient at the application layer. Asymmetric (public key) cryptography is used in support of the digital signature service and encryption key management. Symmetric (secret key) cryptography is used in support of the encryption service. The procedures are intended to be compatible with a wide range of public key management approaches, including both ad hoc and certificate-based schemes. Mechanisms are provided to support many public key management approaches.”

NYT-9-19-95 J. Markoff, “Security Flaw is Discovered in Software Used in Shopping”, the New York Times Front Page, Sept. 19, 1995

NYT-9-20-95 L. Zuckerman, “AT&T Starts On-Line Service Aimed as Small Business”, the New York Times Business Section, Sept. 20, 1995

RFC-1244 “Site Security Handbook” (available from ftp://nis.nsf.net/internet/documents/rfc/rfc1244.txt or ftp://ds.internic.net/rfc/rfc1244.txt)

“This handbook is the product of the Site Security Policy Handbook Working Group (SSPHWG), a combined effort of the Security Area and User Services Area of the Internet Engineering Task Force (IETF).”

“This handbook is a guide to setting computer security policies and procedures for sites that have systems on the Internet. This guide lists issues and factors that a site must consider when setting their own policies. It makes some recommendations and gives discussions of relevant areas.”

RFC-1244 is updated by [SSH].

RSA http://www.rsa.com/rsalabs/faq/faq_home.html

From the introduction to the FAQ:

“This is an introduction to modern cryptography, including answers to commonly asked questions about public key algorithms such as RSA, ElGamal and Diffie-Hellman; secret key techniques such as DES, RC2 and RC4; and hash functions such as MD, MD2, MD5 and SHA. Certificates, key management, patents, Kerberos, discrete log, factoring, domestic and international standards are also among the topics discussed.”

“New in this edition is expanded treatment of recent government involvement in encryption policy and standards, including discussions on the controversial Capstone, Clipper and DSS proposals, export controls, NIST, NSA, privacy and intellectual property concerns.”

SHTTP E. Rescorla, A. Schiffman, “The Secure HyperText Transfer Protocol” (available from ftp://ds.internic.net/internet-drafts/draft-ietf-wts-shttp-00.txt)

"This memo describes a syntax for securing messages sent using the Hypertext Transfer Protocol (HTTP), which forms the basis for the World Wide Web. Secure HTTP (S-HTTP) is an extension of HTTP, providing independently applicable security services for transaction confidentiality, authenticity/integrity and non-repudiability of origin."

Spafford "COAST/Spaf's Hotlist: Computer Security, Law & Privacy" http://www.cs.purdue.edu/homes/spaf/hotlists/csec.html

This is a WWW site with many pointers to security related information:

- Organizations & Agencies
  - FIRST Teams (Forum of Incident Response and Security Teams)
  - Professional Organizations
  - U.S. Government
  - Others
- Education in Computer Security
- Publications
  - Newsletters and Mailing Lists
  - FAQs and Glossaries
  - Books & Book Info
  - Other Publications
- Security Archives, Servers & Indices
  - Comprehensive Sites
  - Tools
  - "Underground" Sites
- Cryptography
  - PGP-related
  - Export Control & Politics
  - Other cryptography
- Computer Viruses
- Privacy Issues
- Computing Ethics
- Security in WWW
- Commercial Sites
  - Computer Vendors
  - Primarily Firewalls
  - Others
- Law
- Miscellaneous

SSH B. Fraser, et al, "Site Security Handbook", an Internet-Draft (such drafts are replaced by RFCs or new drafts after about 6 months.) This draft is available from ftp://ds.internic.net/internet-drafts/draft-ietf-ssh-handbook-00.txt

From the IETF SSH home page (http://www.cert.dfn.de/eng/resource/ietf/ssh/):

"The Site Security Handbook Working Group is chartered to create two documents: (1) a revised handbook that will help system and network administrators develop their own site-specific policies and procedures to deal with computer security problems and their prevention and (2) a new handbook for users. The text of these documents will be developed from the existing RFC 1244, plus needed revisions and additions. "

SSL Netscape Communications Corporation, "Netscape SSLRef" (http://home.netscape.com/info/sslref.html):

"SSL is Secure Sockets Layer, an open, publicly available and license-free security protocol specification suitable for use on the Internet and other TCP/IP networks in a broad range of contexts. It can be used with application-level protocols such as HTTP, FTP, Gopher, Telnet, NNTP, rdist, and many others (including protocols yet to be invented).

The SSL protocol enables advanced security in an application using a variety of mechanisms which include authentication, confidentiality and integrity.

The full protocol specification has been available to software developers and the Internet community since last October.

SSLNews Netscape Communications Corporation, "Industry Leaders Support Secure Sockets Layer for Internet Security" (A "news release". (http://home.netscape.com/info/newsrelease17.html):

"MOUNTAIN VIEW, Calif. (March 20, 1995) -- Netscape Communications Corporation today announced that a number of industry-leading companies and organizations are supporting the Secure Sockets Layer (SSL) protocol for Internet security. Apple Computer, Inc., Bank of America, ConnectSoft, Delphi Internet Services Corporation, Digital Equipment Corporation, First Data Corporation, IBM, MarketNet, MasterCard International Inc., MCI Communications Corp., Microsoft Corporation, Novell, Inc., Open Market, Prodigy, Silicon Graphics, Inc., StarNine, Sun Microsystems, Inc., Visa International, and Wells Fargo are among companies backing SSL.

SSL is an open protocol for securing data communications across computer networks. The broad support for this protocol will promote interoperability between products from many organizations and will speed the growth of electronic commerce on the Internet and private TCP/IP networks. Today, more than 3 million people are already using SSL-enabled products, which have been available since December 1994. In October 1994, Netscape published the specification for SSL on the Internet. Recently, the company also published the source code to the reference implementation, called SSLRef, on the net. SSLRef is free for non-commercial use and is available for flat-fee licensing by companies who want to use it in commercial products."

SSLProtocol Netscape Communications Corporation, "The SSL Protocol", http://home.netscape.com/info/SSL.html

Stallings-95a William Stallings, *Network and Internetwork Security,* Prentice Hall, Englewood Cliffs, New Jersey, 1995. (ISBN: 07803-1107-8)

Stallings-95b William Stallings, *Protect Your Privacy - A Guide for PGP Users*, Prentice Hall, Englewood Cliffs, New Jersey, 1995. (ISBN: 0-13-185596-4)

This book is Stallings' description of Phil Zimmerman's PGP system.

USPS Informal discussions with the Electronic Commerce Services group of the U. S. Postal Service about their electronic commerce services plans, especially for certificates, indicates the following.

USPS is organizing their certificate services around a Certification Authority (CA) that they will run without reference to a higher authority (e.g. the IRPA). Like many others who are running CAs, USPS intends to set its policy in accordance with what they think that their commercial customer base will find most useful. USPS expects to see a number of "high level" CAs (like USPS) cooperate by signing each other's certificates in order to allow for inter-operation. (Steven Kent, IETF PEM Working Group Chair. recently indicated that the IETF would soon set up a working group to review and revise the RFC-1422 CA model.)

USPS will issue two classes, and several subclasses of certificates:

- organizational or corporate certificates would essentially let organizations operate their own CA, referenced to USPS CA policy
- personal - five levels of certification:
  - biometric (encodes a body measurement in a certificate and on a "smart token"-like card)
  - certified (in-person presentation of "official", picture id)
  - basic (written or digital signature, mailed in)
  - proxy (power-of-attorney - based on signed documents of someone else)
  - generic (anonymous)

The certificates are X.509v3 and will (in the future) be able to reference attribute certificates (this provides a way of certifying information not defined in X.509 format).

Certificate access will initially be via X.400 and SMTP e-mail, with WWW access coming soon. USPS does not plan to allow unrestricted public access to the certificate database. In USPS operated X.500 servers the certificates would be in a private directory. Their model is that they will have "listed" and "unlisted" certificates. Listed certificates may be obtained by anyone, but only when requested by name (i.e. you have to know the distinguished name, you cannot "browse" the certificate database). Unlisted certificates can only be obtained from the certified entity (i.e. only the "owner" will distribute unlisted certificates, but they will be signed by the CA).

Certificates will be RSA Public Key based and/or DSS based.

Revocation lists will be available (initially) only by e-mail request.

USPS also anticipates a couple of related services in the near future:

- digital postmark (you send in a crypto hash of the document that you wish postmarked and they timestamp against WWV-GMT, sign with their private key, and return)
- archiving (they will postmark your document, and then archive it for a period of time specified by the customer - this is just a concept, they do not have a concrete tertiary storage plan yet).

The USPS approach would allow those organizations that want to establish a CA a convenient way (and more neutral than the IPRA) of cross-CA operation.

For more information contact Paul Raines (praines@email.usps.gov), Program Manager, Electronic Commerce Services, USPS.